

Quantization via Empirical Divergence Maximization

Michael A. Lexa
amlexa@gmail.com

5 Nov 2011, Revised 30 Apr 2012

Abstract

Empirical divergence maximization (EDM) refers to a recently proposed strategy for estimating f -divergences and likelihood ratio functions. This paper extends the idea to empirical vector quantization where one seeks to empirically derive quantization rules that maximize the Kullback-Leibler divergence between two statistical hypotheses. We analyze the estimator's error convergence rate leveraging Tsybakov's margin condition and show that rates as fast as n^{-1} are possible, where n equals the number of training samples. We also show that the Flynn and Gray algorithm can be used to efficiently compute EDM estimates and show that they can be efficiently and accurately represented by recursive dyadic partitions. The EDM formulation have several advantages. First, the formulation gives access to the tools and results of empirical process theory that quantify the estimator's error convergence rate. Second, the formulation provides a previously unknown derivation for the Flynn and Gray algorithm. Third, the flexibility it affords allows one to avoid a small-cell assumption common in other approaches. Finally, we illustrate the potential use of the method through an example.

1 Introduction

In statistical learning theory, empirical risk minimization is a standard technique whereby classifiers are formed from empirical data [1]. The idea is simple enough: when the underlying probability distributions characterizing the data are unknown, classifiers are found by minimizing an empirical form of the risk (probability of error) over some specified class of classifiers. The technique is well-understood and has been generalized to include various cost criteria and problem settings. In its generalized form, empirical risk minimization is sometimes referred to as M -estimation (the M standing for minimization or maximization) [2].

Recently, Nguyen, Wainwright and Jordan [3] applied M -estimation to the estimation of f -divergences (the Kullback-Leibler (KL) divergence [4] in particular) and to bounded likelihood ratio functions. In this paper, we build on their ideas and develop a method for computing empirical quantization rules by maximizing the KL divergence. We call the method *empirical divergence maximization* (EDM) in deference to its similarity to empirical risk minimization and because the name is simple and descriptive. The proposed formulation leads to an entirely different algorithm for computing the estimators than that employed in [3], and the convergence rates reported here incorporate a margin condition not included in [3] that shows when fast convergence is possible.

As the name suggests, the criterion used in EDM is the KL divergence, a well-known information theoretic quantity that has enjoyed a prominent and long-standing place in both theory and practice. Applications are numerous and range from detection and estimation problems [5–7] to texture retrieval in image databases [8] and from the study of neural coding [9] to linguistic problems [10]. Roughly speaking, the KL divergence quantifies the dissimilarity

of two probability density functions (pdfs) and is therefore often regarded as a “distance”, although it is not a distance metric. Stein’s Lemma [11, p. 309] fundamentally links the divergence to hypothesis testing by relating it to the decay rate of different error probabilities. In fact, the divergence equals the optimal asymptotic error decay rate of a Neyman-Pearson test. Thus increasing the divergence between two statistical hypotheses generally increases their discriminability.

The use of the KL divergence in quantization problems dates back nearly four decades [5]. In that time, various problem settings have been investigated including scalar, vector, and distributed quantization [5, 12–14]. Until recently, however, most results addressing this type of quantization assumed full knowledge of the probability distributions of interest and did not explicitly address empirical designs. Moreover, those works in the quantization literature most closely related to the present paper [14, 15] invoke a small-cell assumption that forces partitions, designed to maximize the divergence, to resemble nearest neighbor partitions even when such partitions cannot well-approximate theoretically optimal partitions (see Fig. 3). Because of its flexibility, however, the EDM approach overcomes this shortcoming.

In [15], Lazebnik and Raginsky study a conceptually similar quantization problem to the one considered here, but the differences between the approaches are substantial. For example, their information loss criterion is a difference of mutual informations, and while related to the KL divergence, this criterion measures a different quantity than the divergence loss studied here. Their work is also placed in a machine learning setting where the data and the quantization values (labels) are jointly distributed and both play integral roles in their information criterion. In this paper, the quantization values play a secondary role in the computation of our estimator.

To formalize the problem, let P and Q be two probability measures defined on the probability space $([0, 1]^d, \mathcal{B})$, where \mathcal{B} denotes the usual Borel σ -algebra and $d \geq 1$. Let p and q denote the density functions of P and Q with respect to Lebesgue measure and assume P and Q are absolutely continuous with respect to one another. Then any quantization rule $\gamma : \mathbb{R}^d \mapsto \{0, \dots, L-1\}$ that operates on a random vector X (distributed according to P or Q) induces the probability mass functions (pmfs), $p(\gamma) = (p_0(\gamma), \dots, p_{L-1}(\gamma))$ and $q(\gamma) = (q_0(\gamma), \dots, q_{L-1}(\gamma))$, where $p_i(\gamma) = P(\gamma(X) = i)$ and similarly for $q_i(\gamma)$. In this context, the KL divergence is defined as

$$D_{KL}(p(\gamma) \| q(\gamma)) := \sum_{i=0}^{L-1} -p_i(\gamma) \log \left(\frac{q_i(\gamma)}{p_i(\gamma)} \right).$$

In EDM, we maximize an empirical form of the KL divergence over some given class of quantization rules. We therefore analyze an estimator of the form

$$\hat{\gamma}_n = \arg \max_{\gamma \in \Gamma} D_n(\gamma),$$

where $D_n(\gamma)$ represents an empirical KL divergence that is defined in Section 2 (the subscript n signifies that it is an empirical quantity that is based on n samples from both P and Q) and where Γ denotes some class of quantization rules. By design EDM estimators constructed rules $\hat{\gamma}_n$ that induce maximally divergent pmfs, thereby best preserving the discriminability of P and Q . In other words, EDM estimators maximize the performance (in terms of KL divergence) of any downstream detector or classifier that operates on the quantized data. The EDM formulation has several advantages: (i) it readily permits the application of empirical process theory which in turn provides the tools to quantify the estimator’s error decay rates; (ii) it naturally leads to the Flynn and Gray algorithm which efficiently computes the quantization

rules; (iii) it provides a systematic derivation for the Flynn and Gray algorithm; and (iv) the flexibility in candidate function classes allows efficient representation of the quantization rules and overcomes the small-cell constraint.

2 Empirical Divergence Maximization

The form of $D_n(\gamma)$ is taken from recent work by Nguyen et al. [3] and relies on rewriting the convex function $-\log(\cdot)$ appearing in the definition of the KL divergence. Throughout the paper, we hold the number of quantization levels L fixed.

2.1 Expressing divergence using convex conjugates

The notion of a *convex conjugate* is based on the observation that a curve can either be described by its graph or by an envelope of tangents [16]. More concretely, a (closed) convex function $f : \mathbb{R} \mapsto \mathbb{R}$ can be described as the pointwise supremum of a collection of affine functions $h(t) = tt^* - \mu^*$ such that the set of all pairs (t^*, μ^*) lie within the epigraph of its convex conjugate $f^*(t^*)$, i.e.,

$$f(t) = \sup_{t^*} \{t^*t - f^*(t^*)\} \quad (1)$$

where by duality the convex conjugate $f^*(t^*)$ of $f(t)$ is defined by

$$f^*(t^*) = \sup_t \{tt^* - f(t)\}.$$

Now suppose γ is an arbitrary quantization rule defined on $[0, 1]^d$,

$$\gamma(x) = \sum_{i=0}^{L-1} i \mathbf{1}_{R_i}(x), \quad x \in [0, 1]^d, \quad (2)$$

where $\{R_i\}_{i=0}^{L-1}$ is a collection of disjoint sets partitioning $[0, 1]^d$ and $\mathbf{1}_{R_i}(\cdot)$ denotes the indicator function. Using (1), we can write the divergence between the pmfs induced by γ as

$$D_{KL}(p(\gamma) \| q(\gamma)) = \sum_{i=0}^{L-1} p_i(\gamma) f\left(\frac{q_i(\gamma)}{p_i(\gamma)}\right) \quad (3a)$$

$$= \sum_{i=0}^{L-1} p_i(\gamma) \cdot \sup_{t^*} \left\{ t^* \frac{q_i(\gamma)}{p_i(\gamma)} - f^*(t^*) \right\}, \quad (3b)$$

where $f(t) = -\log(t)$ for $t > 0$ and $+\infty$ otherwise. Calculating the convex conjugate, one finds

$$f^*(t^*) = \begin{cases} -1 - \log(-t^*) & \text{if } t^* < 0 \\ +\infty & \text{if } t^* \geq 0. \end{cases}$$

Substituting this expression into (3b), we have the following expressions for the KL divergence

$$\begin{aligned}
 D_{KL}(p(\gamma)\|q(\gamma)) &= \sum_{i=0}^{L-1} p_i(\gamma) \sup_{t_i^* \in \mathbb{R}^-} \left\{ t_i^* \frac{q_i(\gamma)}{p_i(\gamma)} + 1 + \log(-t_i^*) \right\} \\
 &= \sum_{i=0}^{L-1} p_i(\gamma) \sup_{c_{R_i} \in \mathbb{R}^+} \left\{ \log(c_{R_i}) - c_{R_i} \frac{q_i(\gamma)}{p_i(\gamma)} + 1 \right\} \\
 &= 1 + \sum_{i=0}^{L-1} \sup_{c_{R_i} \in \mathbb{R}^+} \left\{ P(R_i) \log(c_{R_i}) - c_{R_i} Q(R_i) \right\},
 \end{aligned}$$

where in the second step we let $c_{R_i} = -t_i^*$, and in the last step use the fact that $p_i(\gamma) = P(R_i)$, $R_i = \{x : \gamma(x) = i\}$. The validity of last expression is easily verified by differentiating it with respect to c_{R_i} and solving for the maximizers. By defining the piecewise constant function

$$\phi(x) := \sum_{i=0}^{L-1} c_{R_i} \mathbf{1}_{R_i}(x), \quad c_{R_i} \in \mathbb{R}^+, x \in [0, 1]^d \quad (4)$$

we can write $D_{KL}(p(\gamma)\|q(\gamma))$ in integral form:

$$1 + \sup_{\phi} \left\{ \int_{[0,1]^d} \log(\phi) dP - \int_{[0,1]^d} \phi dQ \right\}, \quad (5)$$

where the supremum is taken over all functions of the form (4). Note that the ϕ which achieves the supremum depends on P , Q , and $\{R_i\}_{i=0}^{L-1}$. Below, we restrict ϕ to lie within a (more) specific class of rules and define the proposed quantization rule estimator in terms of the empirical counterpart to (5).

In addition, note that unlike γ , ϕ does not map $[0, 1]^d$ to a set of indices. We nevertheless refer to both as quantization rules since ϕ only assumes L real values. Note also that in terms of KL divergence, ϕ determines γ , i.e., if ϕ is known, a quantization rule $\gamma : [0, 1]^d \mapsto \{0, \dots, L-1\}$ can be defined that induces the same pmfs as ϕ . This fact becomes important for the algorithm described in Section 3.

2.2 Empirical estimator

To define a function class for ϕ , we first consider different “labelings” of a uniform partition of $[0, 1]^d$. For a given positive integer J , let π_J denote a tessellation of $[0, 1]^d$ by uniform hypercubes $S_k, k = 0, \dots, 2^{dJ} - 1$. To each cell S_k , we can associate one of L labels $\{0, \dots, L-1\}$, and thus for each different labeling of π_J , we can define another partition, π_R , with cells $\{R_i\}_{i=0}^{L-1}$ described by

$$R_i = \bigcup_{k: \text{label}(S_k)=i} S_k, \quad i = 0, \dots, L-1. \quad (6)$$

Now, for a given partition π_R and positive constants $m > 0$ and $M < \infty$, denote by $\Phi_{\pi_R}(L, J, m, M)$ the set of all L -level piecewise constant functions defined on π_R that are bounded and positive:

$$\Phi_{\pi_R}(L, J, m, M) = \left\{ \phi(x) = \sum_{i=0}^{L-1} c_{R_i} \mathbf{1}_{R_i}(x) : m \leq c_{R_i} \leq M \right\}.$$

Letting Π_R denote the set of all partitions π_R , or equivalently the set of all different labelings of π_J , we define the candidate class $\Phi(L, J, m, M)$ of our empirical quantizers as

$$\Phi(L, J, m, M) := \bigcup_{\pi_R \in \Pi_R} \Phi_{\pi_R}(L, J, m, M). \quad (7)$$

Letting $\{X_i^p\}_{i=1}^n$ and $\{X_i^q\}_{i=1}^n$ be training data distributed according to p and q , respectively, we define the function $D_n(\phi)$ as an empirical counterpart to (5)

$$D_n(\phi) := 1 + \frac{1}{n} \sum_{i=1}^n \log \phi(X_i^p) - \frac{1}{n} \sum_{i=1}^n \phi(X_i^q) \quad (8)$$

and define the proposed empirical quantization rule estimator as

$$\hat{\phi}_n := \arg \max_{\phi \in \Phi(L, J, m, M)} D_n(\phi). \quad (9)$$

$\hat{\phi}_n$ is our *empirical divergence maximization* (EDM) estimator. Note $D_n(\phi)$ is not in general a KL divergence; it can in fact be negative for some $\phi \in \Phi$. It is a consistent estimator, however, converging to the “best in class” estimator as $n \rightarrow \infty$ [3].

2.3 Best in class and optimal quantization rules

The *best in class estimate* ϕ^* is that element in Φ that maximizes $D(\phi)$,

$$\begin{aligned} \phi^* &:= \arg \max_{\phi \in \Phi(L, J, m, M)} D(\phi), \quad \text{where} \\ D(\phi) &:= 1 + \int_{[0,1]^d} \log(\phi) dP - \int_{[0,1]^d} \phi dQ. \end{aligned} \quad (10)$$

Note that $D(\phi)$, as opposed to $D_n(\phi)$, is not an empirical quantity; its definition requires full knowledge of the distributions P and Q .

We take the theoretically optimal quantization rule ψ^* to be the rule that maximizes the divergence over a class of piecewise constant functions that has an assumed boundary regularity (the regularity conditions play a role in the convergence analysis in Section 4.2). The class definition uses the notion of a locally constant function: a function $f : [0, 1]^d \mapsto \mathbb{R}$ is *locally constant* at a point $x \in [0, 1]^d$ if there exists $\epsilon > 0$ such that for all $y \in [0, 1]^d$, the condition $\|x - y\| < \epsilon$ implies $f(y) = f(x)$.

Definition (PC class [17]). *A function $f : [0, 1]^d \mapsto \{c_i\}_{i=0}^{L-1}$, $c_i \in \mathbb{R}^+$ is a positive-valued piecewise constant function with L levels if it is locally constant at any point $x \in [0, 1]^d \setminus B(f)$, where $B(f) \subset [0, 1]^d$ is a boundary set satisfying $N(r) \leq \beta r^{-(d-1)}$ for all $r > 0$. Here, $\beta > 0$ is a constant and $N(r)$ is the minimal number of balls of diameter r that covers $B(f)$. Furthermore, let f be uniformly bounded on $[0, 1]^d$, that is $m \leq f(x) \leq M$ for all $x \in [0, 1]^d$, where $m > 0$ and $M < \infty$. The set of all piecewise constant functions f satisfying the above conditions is denoted by $\text{PC}(\beta, m, M, L)$.*

In short, we consider $\text{PC}(\beta, m, M, L)$ to be a class of likelihood-ratio quantization rules that have well behaved boundaries. The theoretically optimal quantization rule is thus defined to be

$$\psi^* := \arg \max_{\psi \in \text{PC}(\beta, m, M, L)} D(\psi). \quad (11)$$

It is well-known that ψ^* can always be constructed by thresholding the likelihood ratio [18]. In other words, the optimal quantization rule ψ^* can always be chosen to be a piecewise constant function whose boundary sets are level sets of the likelihood ratio $q(x)/p(x)$.

3 Solving for the estimator

To find $\hat{\phi}_n$ in (9), we employ a modified form of the Flynn and Gray algorithm [19] that iteratively maximizes the divergence between two pmfs over a set of quantization rules. The method directly follows from the EDM formulation (although it was not originally proposed in this context) and searches for an optimal cell labeling for a given partition where the number of cells is much larger than the number of quantization levels.

3.1 The Flynn and Gray algorithm

For independent and identically distributed random variables X_1, \dots, X_n , the *empirical measure* of a set $A \in [0, 1]^d$, denoted $P_n(A)$, is the sample average

$$P_n(A) = \frac{1}{n} \sum_{k=1}^n \mathbf{1}_A(X_k). \quad (12)$$

The sample average of a function $g : [0, 1]^d \mapsto \mathbb{R}$ can thus be written with respect to P_n as an *empirical expectation*,

$$P_n(g) = \frac{1}{n} \sum_{k=1}^n g(X_k) = \int g dP_n. \quad (13)$$

Using this notation, we rewrite (8) as

$$D_n(\phi) = 1 + \int_{[0,1]^d} \log(\phi) dP_n - \int_{[0,1]^d} \phi dQ_n, \quad (14)$$

where $\phi \in \Phi$. For any fixed partition $\pi_R \in \Pi_R$, $D_n(\phi)$ is maximized by assigning $\phi(x)$ the values $P_n(R_i)/Q_n(R_i)$ for $x \in R_i$. For this assignment choice, $D_n(\phi)$ can be expressed as

$$D_n(\phi) = 1 + \sum_{i=0}^{L-1} \int_{R_i} \log\left(\frac{P_n(R_i)}{Q_n(R_i)}\right) dP_n - \int_{R_i} \frac{P_n(R_i)}{Q_n(R_i)} dQ_n, \quad (15)$$

and the estimator $\hat{\phi}_n$ can now be found by searching over Π_R for the partition that maximizes (15). The Flynn and Gray algorithm [19] is a straightforward method which accomplishes this task. To apply it, we rewrite (15) as

$$D_n(\phi) = \sum_{i=0}^{L-1} P_n(R_i) \left[\log\left(\frac{P_n(R_i)}{Q_n(R_i)}\right) + 1 \right] + Q_n(R_i) \left(-\frac{P_n(R_i)}{Q_n(R_i)} \right) \quad (16)$$

$$= \sum_{i=0}^{L-1} P_n(R_i) a_i + Q_n(R_i) b_i \quad (17)$$

$$= \sum_{i=0}^{L-1} \sum_{k \in I_i} P_n(S_k) a_i + Q_n(S_k) b_i, \quad (18)$$

where $a_i = \log(P_n(R_i)/Q_n(R_i)) + 1$, $b_i = -P_n(R_i)/Q_n(R_i)$, and I_i is the index set $\{k : \text{label}(S_k) = i\}$. The algorithm maximizes $D_n(\phi)$ by iterating two steps: it first holds the set of weights $\{a_i\}$ and $\{b_i\}$ fixed and finds the labels for each cell $S_k \in \pi_J, k = 0, \dots, 2^{dJ} - 1$ that maximizes (18), and then holds the cell labels of π_J fixed and updates the weights $\{a_i\}$ and $\{b_i\}$ using the probabilities $P_n(R_i), Q_n(R_i), i = 0, \dots, L - 1$ found from the first step. Flynn and Gray showed these steps monotonically increase (18), and since $D_n(\phi)$ is upper bounded by $1 - m + \log M$ (follows from the boundedness of ϕ), the algorithm converges to a local maximum. The algorithm returns a locally optimal labeling of π_J and locally optimal weights from which $\hat{\phi}_n$ can be determined:

$$\hat{\phi}_n(x) = -b_i \quad \text{for } x \in R_i. \quad (19)$$

The algorithm is outlined in the panel entitled Algorithm 1. An advantage of the Flynn and Gray algorithm is that it avoids the exhaustive combinatoric search over all possible labelings by only needing to examine each cell S_k once per iteration. From experiments, it has been observed that for moderate sized partitions π_J ($< 2^{16}$ cells) and for $L < 10$, the algorithm converges very quickly (< 30 iterations).

EDM provides a new derivation for the Flynn and Gray algorithm; however, it is interesting to note that it can also be based on the fact that

$$D_{KL}(p||q) = \sup_{\gamma} D_{KL}(p(\gamma)||q(\gamma)) \quad (20)$$

where the supremum is over all measurable quantization rules with an arbitrary number of levels [20, 21]. Because $D_{KL}(p||q) \geq D_{KL}(p(\gamma)||q(\gamma))$ for any quantization rule, one could use (20) to justify an approach similar to EDM and maximize $D_{KL}(p(\gamma)||q(\gamma))$ over a set of quantization rules for a fixed quantization level. While this approach leads to similar (if not identical) estimators, EDM has the advantage of making a clear connection with empirical process theory which provides the theoretical tools to analyze the error convergence rate.

Note also that the original Flynn and Gray algorithm does not explicitly constrain the values of ϕ to lie within the range $[m, M]$. However, to avoid computing unbounded estimates at any given iteration, we employ the K-T technique [22] when computing $P_n(S_k)$ and $Q_n(S_k)$, $k = 0, \dots, 2^{dJ} - 1$. This technique simply preloads each cell S_k by one half before calculating the sample averages, thereby avoiding the possibility of computing zero probability estimates $P_n(R_i), Q_n(R_i)$. Thus instead of (12), one computes

$$P_n(S_k) = \frac{1}{2^{dJ-1} + n} \left(\frac{1}{2} + \sum_{k=1}^n \mathbf{1}_{S_k}(X_k^p) \right) \quad (21)$$

and likewise for $Q_n(S_k)$. Here, the choice of $1/2$ is not arbitrary; it is based on theoretical considerations of what *a priori* distribution of the probabilities $P(S_k), Q(S_k)$ influences the sample averages $P_n(S_k), Q_n(S_k)$ the least [22].

In short, one solves for an EDM estimator based upon the training data $\{X_i^p\}, \{X_i^q\}$ by first computing the sample averages $P_n(S_k), Q_n(S_k)$ for each cell $S_k \in \pi_J$ and then providing these probabilities as input to the Flynn and Gray algorithm. The algorithm is applied to two numerical examples in Section 5.


3.2 Recursive dyadic partitions

Because Φ is based on a uniform dyadic partition, any EDM estimate $\hat{\phi}_n$ can be viewed as a piecewise constant function supported on a *recursive dyadic partition* (RDP). RDPs are a

Algorithm 1 Modified Flynn and Gray algorithm

Input: L, J, n , empirical cell probabilities $P_n(S_k)$
and $Q_n(S_k)$, stopping threshold ϵ

- 1: Initialize iteration index $l = 0$
 - 2: Randomly label cells $S_k \in \pi_J$
 - 3: Compute $P_n^{(l)}(R_i), Q_n^{(l)}(R_i), i = 0, \dots, L - 1$
 - 4: Initialize weights
 $a_i^{(l)} = \log(P_n^{(l)}(R_i)/Q_n^{(l)}(R_i)) + 1$
 $b_i^{(l)} = -P_n^{(l)}(R_i)/Q_n^{(l)}(R_i)$
 - 5: Compute $D_n^{(l)}(\phi)$
 - 6: Initialize intermediary divergence $\tilde{D}_n(\phi) = 0$
 - 7: **while** $(D_n^{(l)}(\phi) - \tilde{D}_n(\phi))/D_n^{(l)}(\phi) > \epsilon$ **do**
 - 8: Find new label for each cell S_k by computing
 $i^{l+1} = \arg \max_{i \in \{0, \dots, L-1\}} P_n(S_k)a_i^{(l)} + Q_n(S_k)b_i^{(l)}$
(hold weights fixed)
 - 9: Update probabilities for $i = 0, \dots, L - 1$
 $P_n^{(l+1)}(R_i) = \sum_{k: \text{label}(S_k)=i} P_n(S_k)$
 $Q_n^{(l+1)}(R_i) = \sum_{k: \text{label}(S_k)=i} Q_n(S_k)$
(includes K-T preloading if necessary)
 - 10: Compute intermediary divergence
 $\tilde{D}_n(\phi) = \sum_{i=0}^{L-1} P_n^{(l+1)}(R_i)a_i^{(l)} + Q_n^{(l+1)}(R_i)b_i^{(l)}$
 - 11: Update weights
 $a_i^{(l+1)} = \log(P_n^{(l+1)}(R_i)/Q_n^{(l+1)}(R_i)) + 1$
 $b_i^{(l+1)} = -P_n^{(l+1)}(R_i)/Q_n^{(l+1)}(R_i)$
 - 12: Compute new divergence
 $D_n^{(l+1)}(\phi) = \sum_{i=0}^{L-1} P_n^{(l+1)}(R_i)a_i^{(l+1)} + Q_n^{(l+1)}(R_i)b_i^{(l+1)}$
 - 13: $l = l + 1$
 - 14: **end while**
- Output:** $\hat{\phi}_n$ (locally optimal labels of π_J and
weights $\{a_i, b_i\}$), $D_n(\hat{\phi}_n)$
-


 Figure 1: An example two-dimensional RDP ($J = 3$).

systematic class of partitions that have proven to be effective in function estimation and classification problems [17, 23]. Their usefulness stems from their ability to adapt to boundaries (including $PC(\beta, m, M, L)$), thus allowing efficient computation of estimators and concise encoding of estimates. In the present context, RDPs are important because they allow efficient encoding of $\hat{\phi}_n$, and their properties are key in the approximation error analysis presented in Section 4.2.

RDPs are partitions composed of quasi-disjoint sets¹ whose union equals the entire space $[0, 1]^d$. A RDP is any partition that can be constructed using only the following rules [17]:

1. $\{[0, 1]^d\}$ is a RDP.
2. Let $\pi = \{S_0, \dots, S_{k-1}\}$ be a RDP, where $S_i = [u_{i1}, v_{i1}] \times \dots \times [u_{id}, v_{id}]$. Then

$$\pi' = \{S_0, \dots, S_{i-1}, S_i^0, \dots, S_i^{(2^d-1)}, S_{i+1}, \dots, S_{k-1}\}$$

is a RDP, where $\{S_i^0, \dots, S_i^{(2^d-1)}\}$ is obtained by dividing the hypercube S_i into 2^d quasi-disjoint hypercubes of equal size. Formally, let $q \in \{0, \dots, 2^d-1\}$ and $q = q_1 q_2 \dots q_d$ by the binary representation of q . Then

$$S_i^{(q)} = \left[u_{i1} + \frac{v_{i1} - u_{i1}}{2} q_1, v_{i1} + \frac{u_{i1} - v_{i1}}{2} (1 - q_1) \right] \times \dots \times \left[u_{id} + \frac{v_{id} - u_{id}}{2} q_d, v_{id} + \frac{u_{id} - v_{id}}{2} (1 - q_d) \right].$$

We say a RDP has maximal depth J if the side length of its smallest hypercube equals 2^{-J} . Fig. 1 illustrates a RDP approximating an elliptical boundary.

It should be clear RDPs describe tree structures where the root node is the entire space $[0, 1]^d$ and the leaf nodes represent the different cells comprising the RDP. Each branch can have different depths and thus the cells can have different sizes. This property allows a RDP to have larger cells in locations where the function value is constant and smaller cells where the values change (around boundaries). The combination of the systematic tree structure and the partition's adaptivity allow the estimator to be efficiently encoded, that is, the number of bits necessary to map an observation to its quantized value can be done efficiently [23].

For a fixed estimator $\hat{\phi}_n$ (or a fixed labeling of π_J), a RDP can be easily constructed by repeating step 2 above (starting with the whole space), but only producing a split if the cells

¹Two sets are quasi-disjoint if and only if their intersection has Lebesgue measure zero.

$S_k \in \pi_J$ on either side of the split, but within the hypercube of interest, have different values (labels).

4 Error Decay Rates

We gauge the quality of $\hat{\phi}_n$ by characterizing the decay rate of the estimation and approximation errors. *Estimation error* is defined as the difference $D(\phi^*) - D(\hat{\phi}_n)$ and quantifies the error caused by computing $\hat{\phi}_n$ without knowledge of p and q . As the number of samples n increases, the estimation error decreases at a rate (exponent of n) that depends on the complexity of Φ and on the properties of p and q . *Approximation error* is defined as $D(\psi^*) - D(\phi^*)$ and arises in cases where $\psi^* \notin \Phi$. To quantify its decay, we think of the candidates rules $\phi \in \Phi$ as being supported on RDPs and the rate of decay in terms of the depth parameter J .

We begin in a standard fashion with two basic inequalities that follow from the definitions of ϕ^* and $\hat{\phi}_n$: $D(\phi^*) - D(\hat{\phi}_n) \geq 0$ and $D_n(\phi^*) - D_n(\hat{\phi}_n) \leq 0$. They imply that the estimation error is upper bounded by a difference of empirical processes

$$\begin{aligned} 0 &\leq D(\phi^*) - D(\hat{\phi}_n) \\ &\leq -[(D_n(\phi^*) - D(\phi^*)) - (D_n(\hat{\phi}_n) - D(\hat{\phi}_n))] \\ &= -(\nu_n(\phi^*) - \nu_n(\hat{\phi}_n))/\sqrt{n}, \end{aligned}$$

where the second inequality results from adding and subtracting $D_n(\phi^*)$ and $D_n(\hat{\phi}_n)$, and where $\nu_n(\gamma) = \sqrt{n}(D_n(\gamma) - D(\gamma))$. Adding the approximation error to both sides of the inequality bounds the total error by the two component errors.

$$\begin{aligned} 0 &\leq \underbrace{D(\psi^*) - D(\hat{\phi}_n)}_{\text{total error}} \\ &\leq \underbrace{-(\nu_n(\phi^*) - \nu_n(\hat{\phi}_n))/\sqrt{n}}_{\text{upper bound on est. error}} + \underbrace{D(\psi^*) - D(\phi^*)}_{\text{approx. error}} \end{aligned} \tag{22}$$

We use this equation and examine the estimation and the approximation errors separately, giving the final rate result for the expected total error.

4.1 Estimation Error

Let \mathbb{E}_p and \mathbb{E}_q denote expectation operator with respect to P and Q . Then, by writing $|\nu_n(\hat{\phi}_n) - \nu_n(\phi^*)|$ as

$$\begin{aligned} &\left| \frac{1}{n} \sum_{i=1}^n (\log \phi(X_i^p) - \log \phi^*(X_i^p)) - \mathbb{E}_p[\log \phi(X) \right. \\ &\quad \left. - \log \phi^*(X)] + (\phi(X_i^q) - \phi^*(X_i^q)) - \mathbb{E}_q[\phi(X) - \phi^*(X)] \right|, \end{aligned}$$

it is clear that for a given quantization rule ϕ , the empirical averages above converge almost surely to their respective values by the strong law of large numbers. But because $\hat{\phi}_n$ can potentially be any element in Φ , any characterization of the convergence rate must hold uniformly over Φ . It is well-known that uniform rates of convergence depend on the complexity of the function class from which the empirical estimators are drawn [24]. Here we use the notion of

bracketing entropy to characterize complexity of Φ . Roughly speaking, the bracketing entropy of a function class \mathcal{G} equals the logarithm of the minimum number of function pairs that upper and lower bound (bracket) all the members in \mathcal{G} to within some tolerance δ and with respect to some norm (a precise definition can be found in [24, p. 16]). We denote the bracketing entropy by $H_B(\delta, \mathcal{G}, L_2(P))$ and say \mathcal{G} has *bracketing complexity* $\alpha > 0$ if $H_B(\delta, \mathcal{G}, L_2(P)) \leq A\delta^{-\alpha}$ for all $\delta > 0$ and for some constant $A > 0$. Because the members of Φ are uniformly bounded, it can be shown that Φ has bracketing complexity $\alpha = 1$ [25]. This fact is incorporated into Theorem 1 below; however, the proof of the theorem given in Appendix A.1 assumes the bracketing complexity lies between zero and two. The proof therefore yields a slightly more general result than that stated.

We now introduce two conditions on p and q . The first simply states that p and q are uniformly bounded.

Condition 1. Assume $c \leq p(x), q(x) \leq C$ for all $x \in [0, 1]^d$, $c > 0$, $C < \infty$.

The second is a condition introduced by Mammen and Tsybakov [26, 27] and involves a key parameter κ that provides insight into when fast convergence rates are possible (i.e., rates faster than $n^{-1/2}$). The condition arises in a slightly different form in function estimation and Bayesian classification problems, and within these contexts, it can be related to the behavior of p and q near a boundary of interest². For example, in Bayesian classification, small κ implies a “steep” regression function at the Bayes decision boundary and thus easier classification; large κ implies a “flat” transition and harder classification. Van de Geer [2] describes κ as an “identifiability” parameter in the sense that it characterizes how well $\phi \in \Phi$ can be distinguished from ψ^* . Because ψ^* is determined by P and Q , this condition is ultimately a condition on these underlying distributions.

Condition 2. There exists constants $K > 0$ and $\kappa \geq 1$ such that for all $\phi \in \Phi$,

$$D(\psi^*) - D(\phi) \geq \|\psi^* - \phi\|_{L_2}^\kappa / K. \quad (23)$$

If κ is small (close to 1) for a given P and Q , the difference $D(\psi^*) - D(\phi)$ is larger for $\phi \in \Phi$ close to ψ^* (those ϕ such that $\|\psi^* - \phi\|_{L_2} \leq 1$) compared to those distributions having larger κ values. Intuitively, this means that such ϕ are more distinguishable from ψ^* for those distributions satisfying Condition 2 with small κ compared to those distributions satisfying Condition 2 with larger κ values, where distinguishability is measured in terms of divergence loss $D(\psi^*) - D(\phi)$ ³. The following result shows that κ effectively characterizes this aspect of the problem, and like for Bayesian classification and estimation, is a key parameter for the error convergence rate.

Theorem 1 (Estimation error). *Let $\hat{\phi}_n$, ϕ^* , and ψ^* be as defined in (9), (10) and (11) respectively. Suppose Conditions 1 and 2 are met for some constants c, C, K , and κ . Then for any $0 < \epsilon < 1$ we have*

$$D(\psi^*) - \mathbb{E}D(\hat{\phi}_n) \leq \left(\frac{1 + \epsilon}{1 - \epsilon} \right) \left[\text{const}(c, C, K, \kappa) n^{-\frac{\kappa}{2\kappa-1}} + D(\psi^*) - D(\phi^*) \right],$$

for sufficiently large n where $\text{const}(c, C, K, \kappa)$ is a decreasing function of ϵ .

²In Bayesian classification and function estimation, the condition is known as the *margin condition*

³Note that the distinguishability in terms of divergence loss is intimately connected with how $\hat{\phi}_n$ is computed: since $D_n(\phi)$ is a surrogate of $D(\phi)$, maximizing $D_n(\phi)$ over $\phi \in \Phi$ is a surrogate for maximizing $D(\phi)$ over $\phi \in \Phi$, or equivalently minimizing $D(\psi^*) - D(\phi)$.




Figure 2: Left: Total expected divergence loss $D(\psi^*) - \mathbb{E}D(\hat{\phi}_n)$ plotted as a function of the number of training samples n (solid curve) for the case $L = 8$, $J = 6$ where $p(x)$ is a zero-mean unit variance Gaussian density and $q(x)$ is a zero-mean unit variance Laplace density. The dashed (black) curve is $\mathcal{O}(n^{-1})$, thus for this example, we have a fast rate of decay. Right: Probability density functions.

The proof is provided in Appendix A.1 (see also [28]) and directly follows from results of van de Geer [2, pp. 206-207] [24] and Mammen and Tsybakov [26].

Depending on P and Q , the decay rate of the estimation error can be as fast as n^{-1} ($\kappa = 1$) and no worse than $n^{-1/2}$ ($\kappa = \infty$). In particular, if for a given P and Q , the approximation error is nonzero (which is commonly the case in quantization problems), we have

$$\begin{aligned} D(\psi^*) - D(\phi) &\geq D(\psi^*) - D(\phi^*) \\ &\geq \text{const} \cdot \frac{\|\psi^* - \phi\|_{L_2}}{M} \end{aligned}$$

where the first inequality follows from the definition of ϕ^* and the second from the fact that $0 \leq \frac{\|\psi^* - \phi\|_{L_2}}{M} \leq 1$ for all $\phi \in \Phi$. Thus with a nonzero approximation error, Condition 2 can be met with $\kappa = 1$ and the rate n^{-1} is achievable. This situation is common in quantization problems because it is unusual in practice for the level sets of a likelihood ratio function to coincide with a RDP for a fixed depth J . For example, consider the simple scenario where $p(x)$ is a zero-mean unit variance Gaussian density function and $q(x)$ is a Laplace density function (also zero-mean and unit variance), $x \in \mathbb{R}$. With $L = 8$ and $J = 6$, the approximation error is 0.002 and hence we expect a rate of $\mathcal{O}(n^{-1})$. Fig. 2 confirms this result experimentally by plotting $D(\psi^*) - \mathbb{E}D(\hat{\phi}_n)$ as a function of n . (For each value of n , Gaussian and Laplacian data were generated and $\hat{\phi}_n$ was computed using the Flynn and Gray algorithm.) The black dashed curve on the left hand plot is shown for reference and equals $40n^{-1} + 0.002$.

In contrast, if P and Q are such that $\psi^* \in \Phi$, then Condition 2 is only met with $\kappa = 2$, and therefore the resulting decay rate is $\mathcal{O}(n^{-2/3})$. To derive this result, consider the subset of quantization rules $\phi \in \Phi$ that share the same partition associated with ψ^* . Recalling (4), we can in this case write ψ^* as

$$\psi^*(x) = \sum_{i=0}^{L-1} \psi_i^* \mathbf{1}_{R_i}(x),$$

where $\{\psi_i^*\}$ are the levels of ψ^* . Then by (10) and the fact that $\psi_i^* = P(R_i)/Q(R_i)$, we have

$$\begin{aligned} D(\psi^*) &= 1 + \int \log(\psi^*) dP - \int \psi^* dQ \\ &= \int \log \psi^* dP \\ &= \int \psi^* \log(\psi^*) dQ \\ &= \sum_{i=0}^{L-1} \int_{R_i} \psi_i^* \log(\psi_i^*) dQ. \end{aligned}$$

Now, because $(\cdot) \log(\cdot)$ is differentiable and continuous on the range of the positive real numbers, we can by Taylor's Theorem [29] expand $\psi_i^* \log \psi_i^*$ on R_i around c_{R_i} for each i to obtain

$$D(\psi^*) = \sum_i \int_{R_i} \psi_i^* \log(c_{R_i}) + \psi_i^* - c_{R_i} + \mathcal{R}_i dQ, \quad (24)$$

where $\mathcal{R}_i = \frac{1}{2\delta_i}(\psi_i^* - c_{R_i})^2$, $i = 0, \dots, L-1$, are the Taylor remainders of the expansions with δ_i lying in between ψ_i^* and c_{R_i} . By adding and subtracting $\int \log(\phi) dP$, (24) can be rearranged to yield

$$\begin{aligned} D(\psi^*) - D(\phi) &= \sum_i \int_{R_i} \mathcal{R}_i dQ \\ &\geq \frac{c}{2M} \sum_i \int_{R_i} (\psi_i^* - c_{R_i})^2 dx \end{aligned} \quad (25)$$

$$= \frac{c}{2M} \|\psi^* - \phi\|_{L_2}^2 \quad (26)$$

where the inequality follows from replacing δ_i with M in the remainder term and using the fact that q is lower bounded by c (Condition 1). Thus, when there is no approximation for the given distributions P and Q (and for given values of J , L , m , and M) the best guaranteed convergence rate is $\mathcal{O}(n^{-2/3})$. Intuitively, this is reasonable since among those ϕ that share the same partition as ψ^* , it is harder to distinguish ψ^* compared to the case where there is a nonzero approximation error.

Nguyen, Wainwright and Jordan reported a similar result to Theorem 1 in [3]. In their investigation, they used an empirical estimator of the same form as (9), but did not consider quantization, nor did they incorporate a margin condition like Condition 2 into their formulation. They considered a class of (inverse) likelihood ratio functions \mathcal{F} that satisfies a complexity condition like Condition 1 and found that the difference $D(f^*) - D_n(\hat{f}_n)$ decays as $\mathcal{O}(n^{-1/(2+\alpha)})$, where $D_n(\cdot)$ and $D(\cdot)$ are as defined in (8) and (10), $f^* \in \mathcal{F}$ is the best in class likelihood ratio function, and \hat{f}_n is an empirical estimator similar to (9). Note that this rate is strictly less than the rate in Theorem 1 even if κ is eliminated from the formulation (take $\kappa \rightarrow \infty$).

4.2 Approximation Error

The approximation error analysis also requires that we now think of Φ as a class of piecewise constant functions (quantization rules) supported on RDPs. As discussed in Section 3.2, this is fully consistent with the definition given in (7). With this in mind, we have the result:

Theorem 2 (Approximation error). *Let $\Phi(L, J, m, M)$, ϕ^* , and ψ^* be as defined in (7), (10), and (11) respectively. Suppose that Condition 1 is met for some constants c and C . Then the approximation error is bounded as*

$$D(\psi^*) - D(\phi^*) \leq \text{const}(\beta, c, C, m, M, L) 2^{-J}. \quad (27)$$

The proof of this result is given in Appendix A.2 (see also [30]). It follows a related, function estimation result in [17] with one important exception: the KL divergence is not additive, thus unlike a mean squared error metric, the approximation error $D(\psi^*) - D(\phi^*)$ cannot be quantified cell by cell. Details are provided in the proof.

The combination of Theorem 1 and Theorem 2 gives the decay rate of the total expected error in terms of the number of training samples n and the depth J of the uniform dyadic partition π_J . To balance the errors and obtain a rate only in terms of n , one can express J as a function of n . Setting $J = \lceil \kappa \ln n / (2\kappa - 1) \ln 2 \rceil$ yields the final result

$$D(\psi^*) - \mathbb{E}D(\hat{\phi}_n) \leq \text{const} \cdot n^{-\frac{\kappa}{2\kappa-1}}, \quad (28)$$

for sufficiently large n .

5 Application: Quantization under communication constraints

When signals are measured and digitized at one location but processed at another, communication of the data is necessary. Because of ever present power, computing, and rate constraints, the raw data cannot be transmitted in full fidelity; instead a summary of the data is sent. When the ultimate goal is classification or detection, one strategy to maximize performance and minimize communication costs is to heavily quantize the data such that the KL divergence is maximized. This is perhaps the simplest strategy and hence attractive when communications are severely constrained. Optimal likelihood-ratio partitions can be very different from typical nearest neighbor (Voronoi) partitions that are associated with quantizers designed to minimize mean squared error (see Figs. 3 and 4). Nevertheless, past work in quantization for classification has forced a small-cell property in the design strategy resulting in partitions resembling nearest neighbor partitions [14]. Consequently, optimal partitions with disjoint regions, for example, cannot be well-approximated by these methods. The EDM quantization method overcomes this shortcoming.

As an illustration, we consider P to be a zero-mean bivariate Gaussian distribution and Q to be a zero-mean bivariate Laplace distribution, both with identity correlation matrices; P and Q thus differ only in their basic shapes. The plot of the likelihood ratio in Fig. 3 shows that the boundaries of the optimal likelihood-ratio partition are concentric circles in each quadrant. Fig. 4(a) depicts the best in class quantization rule along with its associated RDP in Fig. 4(b). The result was generated with the Flynn and Gray algorithm but with $P_n(S_k)$ and $Q_n(S_k)$ in Algorithm 1 replaced by $P(S_k)$ and $Q(S_k)$. (Data points lying outside of $[-5, 5]^2$ were simply ignored.) Convergence occurred in 8 iterations. Fig. 4(c) shows the empirical estimator generated from training sets each of size of two million samples. In this case, the Flynn and Gray algorithm converged in 11 iterations.

In comparison to the best in class quantization rule, Fig. 4 shows the effect of the trying to estimate P and Q on π_J for low probability regions (corner regions). In other words, the lack of data within these regions makes approximating P and Q on π_J difficult, especially by empirical averages. More sophisticated density estimation methods would improve this aspect of the estimator, such as kernel based methods. The estimator might also be improved if one approximates P and Q on a (data-dependent) RDP instead of on π_J (see e.g., [1]).




Figure 3: Plot of a likelihood-ratio function of a zero-mean bivariate Gaussian probability density function and a zero-mean bivariate Laplace probability density function. The level sets of this function, which are concentric circles centered in the quadrants, form the boundaries of an optimal likelihood-ratio partition. Such partitions are not well-approximated by optimal nearest neighbor partitions.

6 Conclusion

In summary, EDM quantization provides a means of finding quantization rules, or more generally, low dimensional transformations that best preserve the divergence between two hypothesized distributions. EDM estimators can be computed using the Flynn and Gray algorithm, and they can exhibit fast error convergence rates as a function of the number of training samples. The EDM formulation benefits from its connection to empirical process theory and possesses the flexibility to overcome the necessity of a small-cell constraint and allow efficient encoding.

Acknowledgments

I thank John Thompson and Mike Davies for their financial support in presenting this work at conferences and for allowing me access to the computational resources of The University of Edinburgh. I also thank Don Johnson for many helpful discussions during the early stages of this work.

A Appendices

A.1 Proof of Theorem 1

The proof proceeds by considering the behavior of $|\nu_n(\hat{\phi}_n) - \nu_n(\phi^*)|$ from (22) as a function of the L_2 distance between $\hat{\phi}_n$ and ϕ^* . This is done by considering a weighted empirical process for two different cases depending on the value of $\|\hat{\phi}_n - \phi^*\|_{L_2}$. The different cases yield different rates of convergence, thus they must be treated separately. At the heart of the argument is Lemma 3, a concentration inequality result by van de Geer [24], Lemma 5.13] concerning the supremum of weighted empirical processes (supremums are considered because we want uniform convergence). The application of this result is not trivial, hence most of the proof is geared toward formulating the problem properly, most of this is done in Case 1 of the proof and Lemma 2. Case 1 is a slight modification of a proof found in [2, pp. 206-207]; Lemma 2 is original. For more information regarding empirical process theory see [24]. Lastly, as stated in Section 4.1, the proof only requires the bracketing complexity of Φ satisfy $0 < \alpha < 2$.

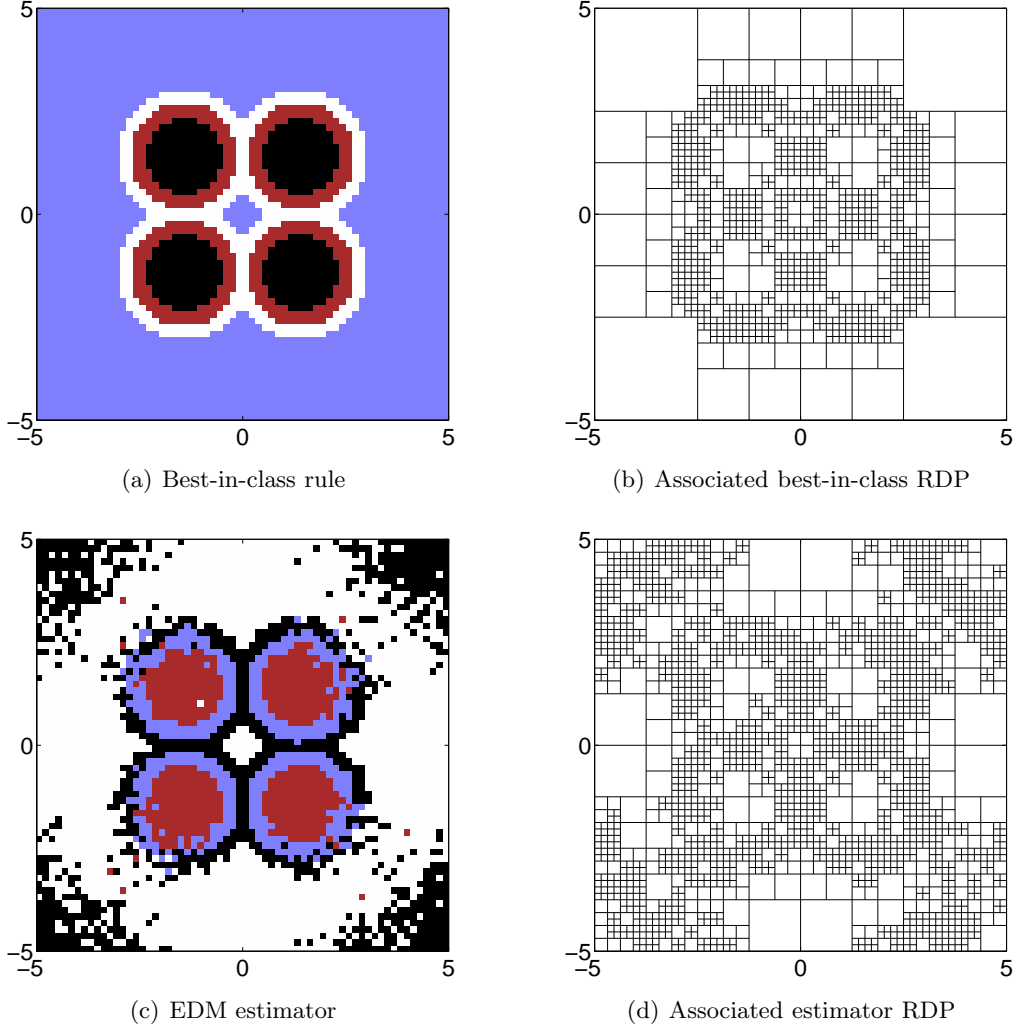


Figure 4: Best-in-class and EDM quantization rules, and their associated recursive dyadic partitions when P is bivariate Gaussian and Q is bivariate Laplace, $L = 4$, $J = 6$. Note that the different cell labelings (colors) are inconsequential in terms of the divergence.

Define the random variable

$$Z_n = \frac{|\nu_n(\widehat{\phi}_n) - \nu_n(\phi^*)|}{C^{(\alpha+2)/4} \|\widehat{\phi}_n - \phi^*\|_{L_2}^{1-\alpha/2} \vee n^{-(2-\alpha)/2(2+\alpha)}}, \quad (29)$$

where $x \vee y = \max(x, y)$. We consider the following two cases: (1) $\|\widehat{\phi}_n - \phi^*\|_{L_2} > C^{-1/2} n^{-1/(2+\alpha)}$ and (2) $\|\widehat{\phi}_n - \phi^*\|_{L_2} \leq C^{-1/2} n^{-1/(2+\alpha)}$.

Case 1. Under this case (29) simplifies to

$$Z_n = \frac{|\nu_n(\widehat{\phi}_n) - \nu_n(\phi^*)|}{C^{(\alpha+2)/4} \|\widehat{\phi}_n - \phi^*\|_{L_2}^\beta} \quad (30)$$

where $\beta = 1 - \alpha/2$. For $\phi = \phi^*$, (30) is defined to be zero. Recalling the inequality (22), we have

$$\begin{aligned} D(\psi^*) - D(\widehat{\phi}_n) &\leq -(\nu_n(\phi^*) - \nu_n(\widehat{\phi}_n))/\sqrt{n} \\ &\quad + D(\psi^*) - D(\phi^*) \\ &\leq Z_n C^{(\alpha+2)/4} \|\widehat{\phi}_n - \phi^*\|_{L_2}^\beta / \sqrt{n} \\ &\quad + D(\psi^*) - D(\phi^*). \end{aligned} \quad (31)$$

Condition 2 implies

$$\begin{aligned} \|\psi^* - \widehat{\phi}_n\|_{L_2}^\beta &\leq K^{\beta/\kappa} (D(\psi^*) - D(\widehat{\phi}_n))^{\beta/\kappa} \\ \|\psi^* - \phi^*\|_{L_2}^\beta &\leq K^{\beta/\kappa} (D(\psi^*) - D(\phi^*))^{\beta/\kappa}. \end{aligned} \quad (32)$$

Hence, by the triangle inequality and (32), we have

$$\begin{aligned} \|\widehat{\phi}_n - \phi^*\|_{L_2}^\beta &\leq \|\psi^* - \widehat{\phi}_n\|_{L_2}^\beta + \|\psi^* - \phi^*\|_{L_2}^\beta \\ &\leq K^{\beta/\kappa} (D(\psi^*) - D(\widehat{\phi}_n))^{\beta/\kappa} \\ &\quad + K^{\beta/\kappa} (D(\psi^*) - D(\phi^*))^{\beta/\kappa}. \end{aligned} \quad (33)$$

Using (33) in (31), shows $D(\psi^*) - D(\widehat{\phi}_n)$ is less than or equal to

$$\begin{aligned} &\left[C^{(\alpha+2)/4} K^{\beta/\kappa} n^{-1/2} Z_n (D(\psi^*) - D(\widehat{\phi}_n))^{\beta/\kappa} + C^{(\alpha+2)/4} \right. \\ &\quad \left. K^{\beta/\kappa} n^{-1/2} Z_n (D(\psi^*) - D(\phi^*))^{\beta/\kappa} \right] + D(\psi^*) - D(\phi^*). \end{aligned}$$

We now apply Lemma 1 to each of the terms within the brackets to obtain

$$\begin{aligned} D(\psi^*) - D(\widehat{\phi}_n) &\leq \epsilon \left[(D(\psi^*) - D(\widehat{\phi}_n)) + (D(\psi^*) - D(\phi^*)) \right] \\ &\quad + 2C^{\frac{-\kappa\beta}{2(\kappa+\beta)}} \left(\frac{K}{\epsilon} \right)^{\frac{\beta}{\kappa-\beta}} Z_n^{\frac{\kappa}{\kappa-\beta}} n^{-\frac{\kappa}{2(\kappa-\beta)}} + D(\psi^*) - D(\phi^*). \end{aligned}$$

By rearranging the previous expression and dropping a factor of $\frac{1}{1+\epsilon} < 1$, we have

$$\begin{aligned} D(\psi^*) - D(\widehat{\phi}_n) &\leq \left(\frac{1+\epsilon}{1-\epsilon} \right) \\ &\quad \left[2C^{\frac{-\kappa\beta}{2(\kappa+\beta)}} (K/\epsilon)^{\frac{\beta}{\kappa-\beta}} Z_n^{\frac{\kappa}{\kappa-\beta}} n^{-\frac{\kappa}{2(\kappa-\beta)}} + D(\psi^*) - D(\phi^*) \right]. \end{aligned} \quad (34)$$

For any $r \geq \frac{\kappa}{\kappa-\beta}$, we have by Jensen's inequality [11], and Lemma 2 that

$$\mathbb{E}Z_n^{\frac{\kappa}{\kappa-\beta}} = \mathbb{E}\left[(Z_n^r)^{\frac{\kappa}{r(\kappa-\beta)}}\right] \leq [\mathbb{E}Z_n^r]^{\frac{\kappa}{r(\kappa-\beta)}} \leq c_2^{\frac{\kappa}{\kappa-\beta}}. \quad (35)$$

Taking the expectation of (34) and applying (35), we conclude

$$\begin{aligned} D(\psi^*) - \mathbb{E}D(\hat{\phi}_n) &\leq \left(\frac{1+\epsilon}{1-\epsilon}\right) \\ &\quad \left[2C^{\frac{-\kappa\beta}{2(\kappa+\beta)}}(K/\epsilon)^{\frac{\beta}{\kappa-\beta}}c_2^{\frac{\kappa}{\kappa-\beta}}n^{-\frac{\kappa}{2(\kappa-\beta)}} + D(\psi^*) - D(\phi^*)\right], \end{aligned}$$

for $\|\hat{\phi}_n - \phi^*\|_{L_2} > C^{-1/2}n^{-1/(2+\alpha)}$.

Case 2. For this case,

$$Z_n = \frac{|\nu_n(\hat{\phi}_n) - \nu_n(\phi^*)|}{n^{-(2-\alpha)/2(2+\alpha)}}.$$

From the fundamental inequality

$$\begin{aligned} D(\psi^*) - D(\hat{\phi}_n) &\leq -(\nu_n(\phi^*) - \nu_n(\hat{\phi}_n))/\sqrt{n} \\ &\quad + D(\psi^*) - D(\phi^*) \\ &\leq Z_n n^{-(2-\alpha)/2(2+\alpha)} n^{-1/2} \\ &\quad + D(\psi^*) - D(\phi^*) \\ &= Z_n n^{-2/(2+\alpha)} + D(\psi^*) - D(\phi^*). \end{aligned} \quad (36)$$

Taking the expectation of (36) and applying Lemma 2 yields

$$D(\psi^*) - \mathbb{E}D(\hat{\phi}_n) \leq c_3 n^{-2/(2+\alpha)} + D(\psi^*) - D(\phi^*), \quad (37)$$

for $\|\hat{\phi}_n - \phi^*\|_{L_2} \leq n^{-1/(2+\alpha)}$. The rate attained in (37) is (strictly) faster than that attained in Case 1 ($n^{-2/(2+\alpha)} < n^{-\kappa/(2(\kappa-1)+\alpha)}$ for $\kappa > 1 - \alpha/2$, $0 < \alpha < 2$). Therefore, the decay of the total divergence loss is governed by the slower rate found in Case 1. ■

Lemma 1 (Tsybakov and van de Geer [31], van de Geer [2]). *We have for all positive v, t , and ϵ , and $\kappa > \beta$,*

$$vt^{\beta/\kappa} \leq \epsilon t + v^{\frac{\kappa}{\kappa-\beta}} \epsilon^{-\frac{\beta}{\kappa-\beta}}.$$

Lemma 2. *Let $\hat{\phi}_n$, ϕ^* , and ψ^* be as defined in (9), (10), and (11) respectively. Then under Conditions 1 and 2, we have*

$$\mathbb{E} \left(\sup_{\phi \in \Phi: \|\phi - \phi^*\|_{L_2} > C^{-1/2}n^{-1/(2+\alpha)}} \frac{|\nu_n(\phi) - \nu_n(\phi^*)|}{\|\phi - \phi^*\|_{L_2}^{1-\alpha/2}} \right)^r \leq c_2^r, \quad (38)$$

and

$$\mathbb{E} \left(\sup_{\phi \in \Phi: \|\phi - \phi^*\|_{L_2} \leq C^{-1/2}n^{-1/(2+\alpha)}} \frac{|\nu_n(\phi) - \nu_n(\phi^*)|}{n^{-(2-\alpha)/2(2+\alpha)}} \right) \leq c_3, \quad (39)$$

for some positive constants c_2, r , and c_3 .

Proof. Case 1: Equation (38). To compact notation, let \mathfrak{z} denote the inequality $\|\phi - \phi^*\|_{L_2} > C^{-1/2} n^{-1/(2+\alpha)}$ and recall that

$$|\nu_n(\phi) - \nu_n(\phi^*)| = \sqrt{n} \left| \int (\log \phi - \log \phi^*) d(P_n - P) + \int (\phi - \phi^*) d(Q_n - Q) \right|.$$

By Condition 1, we have

$$\begin{aligned} \|\phi - \phi^*\|_{L_2(P)} &\leq C^{1/2} \|\phi - \phi^*\|_{L_2} \\ \|\phi - \phi^*\|_{L_2(Q)} &\leq C^{1/2} \|\phi - \phi^*\|_{L_2}, \end{aligned} \tag{40}$$

for $\phi \in \Phi$. Consequently, we can write

$$\begin{aligned} &\sup_{\phi \in \Phi: \mathfrak{z}} \frac{|\nu_n(\phi) - \nu_n(\phi^*)|}{\|\phi - \phi^*\|_{L_2}^{1-\alpha/2}} \\ &\leq \sup_{\phi \in \Phi: \mathfrak{z}} \frac{\sqrt{n} \left| \int (\log \phi - \log \phi^*) d(P_n - P) \right|}{\|\phi - \phi^*\|_{L_2}^{1-\alpha/2}} \\ &\quad + \sup_{\phi \in \Phi: \mathfrak{z}} \frac{\sqrt{n} \left| \int (\phi - \phi^*) d(Q_n - Q) \right|}{\|\phi - \phi^*\|_{L_2}^{1-\alpha/2}} \\ &\leq \sup_{\phi \in \Phi: \mathfrak{z}} \frac{\sqrt{n} \left| \int (\log \phi - \log \phi^*) d(P_n - P) \right|}{C^{(\alpha-2)/4} \|\phi - \phi^*\|_{L_2(P)}^{1-\alpha/2}} \\ &\quad + \sup_{\phi \in \Phi: \mathfrak{z}} \frac{\sqrt{n} \left| \int (\phi - \phi^*) d(Q_n - Q) \right|}{C^{(\alpha-2)/4} \|\phi - \phi^*\|_{L_2(Q)}^{1-\alpha/2}} \end{aligned} \tag{41}$$

where the last inequality follows from (40).

We now want to apply a probability inequality due to van de Geer [24] (stated as Lemma 3 below) to the two terms in (41). The result requires Φ and $\tilde{\Phi} = \{\log \phi : \phi \in \Phi\}$ to have a bracketing complexity satisfying $0 < \alpha < 2$. Φ satisfies the requirement by construction which implies the same is true for $\tilde{\Phi}$. The result also requires that the differences $(\log \phi - \log \phi^*)$ and $(\phi - \phi^*)$ are upper bounded. This follows from the definition of Φ . Furthermore, note that the proper form of the condition under the supremum follows from (40).

Applying Lemma 3 to each term in (41), we obtain

$$\begin{aligned} &\Pr \left(\sup_{\phi \in \Phi: \mathfrak{z}} \frac{\sqrt{n} \left| \int (\log \phi - \log \phi^*) d(P_n - P) \right|}{\|\phi - \phi^*\|_{L_2(P)}^{1-\alpha/2}} \geq C^{(\alpha-2)/4} t \right) \\ &\leq \tilde{c} \exp \left(-\frac{t}{c^2} \right) \end{aligned}$$

and

$$\begin{aligned} \Pr\left(\sup_{\phi \in \Phi: \mathfrak{I}} \frac{\sqrt{n} \left| \int (\phi - \phi^*) d(Q_n - Q) \right|}{\|\phi - \phi^*\|_{L_2(Q)}^{1-\alpha/2}} \geq C^{(\alpha-2)/4} t\right) \\ \leq \tilde{c} \exp\left(-\frac{t}{c^2}\right) \end{aligned}$$

for all $t \geq c$, some constant $\tilde{c} > 0$, and n sufficiently large. Consequently,

$$\mathbb{E}\left(\sup_{\phi \in \Phi: \mathfrak{I}} \frac{\sqrt{n} \left| \int (\log \phi - \log \phi^*) d(P_n - P) \right|}{\|\phi - \phi^*\|_{L_2(P)}^{1-\alpha/2}}\right)^r \leq c_{2,1}$$

and

$$\mathbb{E}\left(\sup_{\phi \in \Phi: \mathfrak{I}} \frac{\sqrt{n} \left| \int (\phi - \phi^*) d(Q_n - Q) \right|}{\|\phi - \phi^*\|_{L_2(Q)}^{1-\alpha/2}}\right)^r \leq c_{2,2}$$

for all $r > 0$ and some finite positive constants $c_{2,1}$ and $c_{2,2}$. Therefore

$$\mathbb{E}\left(\sup_{\phi \in \Phi: \mathfrak{I}} \frac{|\nu_n(\phi) - \nu_n(\phi^*)|}{\|\phi - \phi^*\|_{L_2}^{1-\alpha/2}}\right)^r \leq c_2^r,$$

for some finite positive constant c_2 .

Case 2: Equation (39). Let \mathfrak{b} denote the inequality $\|\phi - \phi^*\|_{L_2} \leq C^{-1/2} n^{-1/(2+\alpha)}$. From the definition of the empirical process $|\nu_n(\phi) - \nu_n(\phi^*)|$, we have

$$\begin{aligned} \sup_{\phi \in \Phi: \mathfrak{b}} \frac{|\nu_n(\phi) - \nu_n(\phi^*)|}{n^{-(2-\alpha)/2(2+\alpha)}} \\ \leq \sup_{\phi \in \Phi: \mathfrak{b}} \frac{\sqrt{n} \left| \int (\log \phi - \log \phi^*) d(P_n - P) \right|}{n^{-(2-\alpha)/2(2+\alpha)}} \\ + \sup_{\phi \in \Phi: \mathfrak{b}} \frac{\sqrt{n} \left| \int (\phi - \phi^*) d(Q_n - Q) \right|}{n^{-(2-\alpha)/2(2+\alpha)}}. \end{aligned} \tag{42}$$

Apply Lemma 3 to each of the terms in (42) to get

$$\begin{aligned} \Pr\left(\sup_{\phi \in \Phi: \mathfrak{b}} \left| \int (\log \phi - \log \phi^*) d(P_n - P) \right| \geq t n^{-2/(2+\alpha)}\right) \\ \leq c \exp\left(-\frac{t n^{\frac{\alpha}{2+\alpha}}}{c^2}\right) \end{aligned}$$

and

$$\begin{aligned} \Pr \left(\sup_{\phi \in \Phi: \flat} \left| \int (\phi - \phi^*) d(Q_n - Q) \right| \geq t n^{-2/(2+\alpha)} \right) \\ \leq c \exp \left(-\frac{t n^{\frac{\alpha}{2+\alpha}}}{c^2} \right), \end{aligned}$$

for all $t \geq c$ and for n sufficiently large. Therefore, for n sufficiently large

$$\mathbb{E} \left(\sup_{\phi \in \Phi: \flat} \frac{|\nu_n(\phi) - \nu_n(\phi^*)|}{n^{-(2-\alpha)/2(2+\alpha)}} \right) \leq c_3,$$

for some positive constant c_3 . □

Lemma 3 (van de Geer [24], Lemma 5.13). *Let X_1, \dots, X_n be an independent and identically distributed sequence of random variables on a probability space $(\mathcal{X}, \mathcal{A}, P)$. Let $\mathcal{G} \subset L_2(P)$ be a collection of functions and define the empirical process indexed by \mathcal{G} as*

$$\nu_n = \left\{ \nu_n(g) = \sqrt{n} \int g d(P_n - P) : g \in \mathcal{G} \right\}.$$

Let $|g|_\infty = \sup_{x \in \mathcal{X}} |g(x)|$ denote the supremum norm and suppose $\sup_{g \in \mathcal{G}} |g - g_0|_\infty \leq K$, for some fixed element $g_0 \in \mathcal{G}$ and some constant K . Furthermore, suppose

$$H_B(\delta, \mathcal{G}, L_2(P)) \leq A\delta^{-\rho}, \quad \text{for all } \delta > 0,$$

for some $0 < \rho < 2$ and some constant $A > 0$. Then for some constant c depending on ρ and A , we have for all $t \geq c$ and for n sufficiently large,

$$\begin{aligned} \Pr \left(\sup_{g \in \mathcal{G}, \|g - g_0\| \leq n^{-\frac{1}{2+\rho}}} \left| \int (g - g_0) d(P_n - P) \right| \geq t n^{-\frac{2}{2+\rho}} \right) \\ \leq c \exp \left(-\frac{t n^{\frac{\rho}{2+\rho}}}{c^2} \right) \end{aligned}$$

and

$$\begin{aligned} \Pr \left(\sup_{g \in \mathcal{G}, \|g - g_0\| > n^{-\frac{1}{2+\rho}}} \frac{|\nu_n(g) - \nu_n(g_0)|}{\|g - g_0\|^{1-\frac{\rho}{2}}} \geq t \right) \\ \leq c \exp \left(-\frac{t}{c^2} \right), \end{aligned}$$

where the norms $\|g - g_0\|$ are norms in $L_2(P)$.

A.2 Proof of Theorem 2

Recall the definition of $\text{PC}(\beta, m, M, L)$ from Section 4.2. We will need the following lemma.

Lemma 4 ([17], Lemma 5, p. 121). *There is a RDP such that the cells intersecting $B(\psi^*)$ are at depth J and all the other cells are at depths no greater than J . Denote the smallest such RDP by π_J^* . Then π_J^* has at most $2^{2d}\beta 2^{(d-1)J}$ cells intersecting $B(\psi^*)$.*

Let ϕ' denote the L -level piecewise constant function defined by,

$$\phi'(x) = \sum_{i=0}^{L-1} c_{R'_i} \mathbf{1}_{R'_i}(x), \quad c_{R'_i} = \frac{P(R'_i)}{Q(R'_i)} \quad (43)$$

where each member region R'_i of the partition $\{R'_i\}$ is composed of a union of cells $S \in \pi_J^*$. Furthermore, let ϕ' satisfy the condition that the cells $S \in \pi_J^*$ contained in $R'_i/B(\psi^*)$ are also contained in A_i^* . More concisely, we write $S \subseteq R'_i/B(\psi^*) \Rightarrow S \subseteq A_i^*/B(\psi^*)$. In words, this last condition means that the partitions $\{R'_i\}$ and $\{A_i^*\}$ coincide except possibly on the boundary $B(\psi^*)$.

First, observe that $D(\psi^*) - D(\phi^*) \leq D(\psi^*) - D(\phi')$ since the divergence between the pmfs induced by ϕ' is necessarily less than or equal to the that induced by the best in class quantization rule ϕ^* . (This inequality also follows from the Data Processing Theorem [4, pp. 18-22].)

Next, upper bound the difference $D(\psi^*) - D(\phi')$ by the L_1 -norm of $(\psi^* - \phi')$:

$$\begin{aligned} D(\psi^*) - D(\phi') &= \int_{[0,1]^d} \log \frac{\psi^*}{\phi'} dP - \int_{[0,1]^d} (\psi^* - \phi') dQ \\ &\leq \int \left(\frac{\psi^*}{\phi'} - 1 \right) dP - \int (\psi^* - \phi') dQ \\ &= \int \frac{1}{\phi'} (\psi^* - \phi') dP - \int (\psi^* - \phi') dQ \\ &\leq \frac{C}{m} \left| \int (\psi^* - \phi') dx \right| + c \left| \int (\psi^* - \phi') dx \right| \\ &\leq \frac{C + cm}{m} \|\psi^* - \phi'\|_{L_1}, \end{aligned} \quad (44)$$

where the first inequality follows from the fact that $\log x \leq x - 1$, for $x > 0$, and the second inequality follows from the bounds on ϕ , p , and q .

Rewrite the L_1 -norm as

$$\begin{aligned} \|\psi^* - \phi'\|_{L_1} &= \int_{[0,1]^d} |\psi^*(x) - \phi'(x)| dx \\ &= \sum_{i=0}^{L-1} \int_{R'_i} |\psi^*(x) - \phi'(x)| dx \\ &= \sum_{i=0}^{L-1} \left[\sum_{S \subseteq R'_i/B(\psi^*)} \int_S |\psi^*(x) - \phi'(x)| dx \right. \\ &\quad \left. + \sum_{S \subseteq R'_i(B(\psi^*))} \int_S |\psi^*(x) - \phi'(x)| dx \right] \end{aligned} \quad (45)$$

Here, $S \subseteq R'_i/B(\psi^*)$ means all cells S that are a subset of R'_i which do not intersect the boundary $B(\psi^*)$. Similarly, $S \subseteq R'_i(B(\psi^*))$ means all cells S that are subsets of R'_i which do intersect $B(\psi^*)$.

Consider the second summation within the brackets in (45). By the boundedness assump-

tions on ψ^* and ϕ' , the integrand can be upper bounded by M . Therefore,

$$\begin{aligned}
 \sum_{S \subseteq R'_i(B(\psi^*))} \int_S |\psi^*(x) - \phi'(x)| dx \\
 \leq M \sum_{S \subseteq R'_i(B(\psi^*))} \text{Vol}(S) \\
 \leq M 2^{2d} \beta 2^{(d-1)J} 2^{-dJ} \\
 = M 2^{2d} \beta 2^{-J},
 \end{aligned} \tag{46}$$

where the second inequality follows from Lemma 4 and the fact that the volume of one cell S is 2^{-dJ} .

Now, consider the first summation within the brackets in (45). For all $S \subseteq R'_i$ (and in particular for all $S \subseteq R'_i/B(\psi^*)$), ϕ' equals $Q(R'_i)/P(R'_i)$ (recall (43)). Likewise, by the definition of ϕ' , ψ^* is also constant for all $S \subseteq R'_i/B(\psi^*)$. Therefore, we have

$$\begin{aligned}
 \sum_{S \subseteq R'_i/B(\psi^*)} \int_S |\psi^*(x) - \phi'(x)| dx \\
 = \sum_{i=0}^{L-1} \left| \frac{P(A_i^*)}{Q(A_i^*)} - \frac{P(R'_i)}{Q(R'_i)} \right| \sum_{S \subseteq R'_i/B(\psi^*)} \text{Vol}(S) \\
 \leq \sum_{i=0}^{L-1} \left| \frac{P(A_i^*)Q(R'_i) - P(R'_i)Q(A_i^*)}{Q(A_i^*)Q(R'_i)} \right| \text{Vol}(R'_i) \\
 \leq \frac{1}{c} \sum_{i=0}^{L-1} \left| \frac{P(A_i^*)Q(R'_i) - P(R'_i)Q(A_i^*)}{Q(A_i^*)} \right|.
 \end{aligned} \tag{47}$$

Using the inequalities,

$$\begin{aligned}
 Q(R'_i) &\leq Q(A_i^*) + \sum_{S \in R'_i(B(A_i^*))} Q(S) \\
 P(R'_i) &\geq P(A_i^*) - \sum_{S \in R'_i(B(A_i^*))} P(S),
 \end{aligned} \tag{48}$$

we upper bound each term in the summation in (47)

$$\begin{aligned}
 \frac{1}{Q(A_i^*)} |P(A_i^*)Q(R'_i) - P(R'_i)Q(A_i^*)| \\
 \leq \frac{1}{Q(A_i^*)} \left| P(A_i^*) \sum_{S \subseteq R'_i(B(A_i^*))} Q(S) \right. \\
 \left. + Q(A_i^*) \sum_{S \subseteq R'_i(B(A_i^*))} P(S) \right| \\
 \leq \frac{1}{Q(A_i^*)} |P(A_i^*)C\beta'2^{-J} + Q(A_i^*)C\beta'2^{-J}| \\
 = C\beta'2^{-J} \left| \frac{P(A_i^*) + Q(A_i^*)}{Q(A_i^*)} \right|,
 \end{aligned}$$

where the second inequality follows from Lemma 4 with $\beta' = 2^{2d}\beta$.

Summarizing, we have

$$\begin{aligned}
& \sum_{i=0}^{L-1} \sum_{S \subseteq R'_i/B(\psi^*)} \int_S |\psi^*(x) - \phi'(x)| dx \\
& \leq \frac{C}{c} \beta' 2^{-J} \sum_{i=0}^{L-1} \left| \frac{P(A_i^*) + Q(A_i^*)}{Q(A_i^*)} \right| \\
& \leq \frac{C}{c} \left(\frac{C}{c} + 1 \right) \beta' L 2^{-J},
\end{aligned} \tag{49}$$

where the last step follows from the assumed bounds on p and q .

Finally, by combining (44), (45), (46), and (49), we conclude

$$\|\psi^* - \phi'\|_{L_1} \leq \beta' L [M + (C/c)(C/c + 1)] 2^{-J} \tag{50}$$

and

$$\begin{aligned}
D(\psi^*) - D(\phi') & \leq \\
& \left(\frac{C + cm}{m} \right) \beta' L [M + (C/c)(C/c + 1)] 2^{-J}
\end{aligned} \tag{51}$$

■

References

- [1] L. Devroye, L. Györfi, and G. Lugosi, *A Probabilistic Theory of Pattern Recognition*, Springer, New York, 1996.
- [2] S. van de Geer, *Lectures on Empirical Processes: Theory and Statistical Applications*, chapter Oracle Inequalities and Regularization, pp. 191–252, European Mathematical Society, 2007.
- [3] XuanLong Nguyen, Martin J. Wainwright, and Michael I. Jordan, “Estimating divergence functionals and the likelihood ratio by convex risk minimization,” *IEEE Trans. Info. Th.*, vol. 56, no. 11, pp. 5847–5861, Nov 2010.
- [4] S. Kullback, *Information Theory and Statistics*, Wiley, New York, 1959.
- [5] H. V. Poor and J. B. Thomas, “Applications of Ali-Silvey distance measures in the design of generalized quantizers for binary decision systems,” *IEEE Trans. on Communications*, vol. 25, no. 9, pp. 893–900, Sep 1977.
- [6] J.F. Chamberland and V.V. Veeravalli, “Decentralized detection in sensor networks,” *IEEE Trans. Signal Processing*, vol. 51, no. 2, pp. 407–416, Feb 2003.
- [7] A.T. Ihler, J.W. Fisher, and A.S. Willsky, “Nonparametric hypothesis tests for statistical dependency,” *IEEE Trans. Signal Processing*, vol. 52, no. 8, pp. 2234–2249, Aug 2004.
- [8] M.N. Do and M. Vetterli, “Wavelet-based texture retrieval using generalized gaussian density and Kullback-Leibler distance,” *IEEE Trans. Image Processing*, vol. 11, no. 2, pp. 146–158, Feb 2002.

- [9] D.H. Johnson, C.M. Gruner, K. Baggerly, and C. Seshargiri, “Information-theoretic analysis of neural coding,” *J. Computational Neuroscience*, , no. 10, pp. 47–69, 2001.
- [10] H. Cai, S. R. Kulkarni, and S. Verdú, “Universal divergence estimation for finite-alphabet sources,” *IEEE Trans. Info. Th.*, vol. 52, no. 8, pp. 3456–3475, Aug 2006.
- [11] T.M. Cover and J.A. Thomas, *Elements of Information Theory*, Wiley, 1991.
- [12] S.A. Kassam, *Signal Detection in Non-Gaussian Noise*, Springer-Verlag, 1988.
- [13] M. Longo, T. D. Lookabaugh, and R. M. Gray, “Quantization for decentralized hypothesis testing under communication constraints,” *IEEE Trans. Info. Th.*, vol. 36, no. 2, pp. 241–255, Mar 1990.
- [14] R. Gupta and A. O. Hero, “High-rate vector quantization for detection,” *IEEE Trans. Info. Th.*, vol. 49, no. 8, pp. 1951–1969, Aug 2003.
- [15] S. Lazebnik and M. Raginsky, “Supervised learning of quantizer codebooks by information loss minimization,” *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 31, no. 7, pp. 1294–1309, Jul 2009.
- [16] R. Tyrrell Rockafellar, *Convex Analysis*, Princeton University Press, Princeton, New Jersey, 1970.
- [17] Rui Castro, *Active Learning and Adaptive Sampling for Non-parametric Inference*, Ph.D., Rice University, Houston TX, U.S.A., Aug. 2007.
- [18] J.N. Tsitsiklis, “Extremal properties of likelihood-ratio quantizers,” *IEEE Trans. Comm.*, vol. 41, no. 4, pp. 550–558, Apr 1993.
- [19] T. Flynn and R. M. Gray, “Encoding of correlated observations,” *IEEE Trans. Info. Th.*, vol. 33, no. 6, pp. 773–787, Nov 1987.
- [20] R.M. Gray, *Entropy and Information Theory*, Springer-Verlag, New York, 1990.
- [21] H.V. Poor, “Fine quantization in signal detection and estimation,” *IEEE Trans. Info. Th.*, vol. 34, no. 5, pp. 960–972, Sep 1988.
- [22] R. E. Krichevsky and V. K. Trofimov, “The performance of universal encoding,” *IEEE Trans. Info. Th.*, vol. 27, no. 2, pp. 199–207, Mar 1981.
- [23] C. Scott and R.D. Nowak, “Minimax-optimal classification with dyadic decision trees,” *IEEE Trans. Info. Th.*, vol. 52, no. 4, pp. 1335 – 1353, April 2006.
- [24] S. van de Geer, *Empirical Processes in M-Estimation*, Cambridge University Press, 2000.
- [25] Michael A Lexa, *Sequential Quantization for Classification: The Impact of Structure and Nonparametric Estimates*, Ph.D., Rice University, Houston TX, U.S.A., Aug. 2008.
- [26] E. Mammen and A. B. Tsybakov, “Smooth discrimination analysis,” *Ann. Statist.*, vol. 27, no. 6, pp. 1808–1829, 1999.
- [27] A. B. Tsybakov, “Optimal aggregation of classifiers in statistical learning,” *Ann. Statist.*, vol. 32, no. 1, pp. 135–166, 2004.

- [28] M.A. Lexa, “Empirical quantization for sparse sampling systems,” *Proc. IEEE Inter. Conf. on Acoustics, Speech, and Signal Processing*, pp. 3942–3945, Mar 2010.
- [29] Watson Fulks, *Advanced Calculus: An Introduction to Analysis*, John Wiley & Sons, New York, 3 edition, 1978.
- [30] M.A. Lexa, “Empirical divergence maximization for quantizer design: An analysis of approximation error,” *Proc. IEEE Inter. Conf. on Acoustics, Speech, and Signal Processing*, pp. 4221–4223, May 2011.
- [31] A. B. Tsybakov and S. A. van de Geer, “Square root penalty: Adaptation to the margin in classification and in edge estimation,” *Ann. Statist.*, vol. 33, no. 3, pp. 1203–1224, 2005.